PCI-SIG ENGINEERING CHANGE NOTICE

TITLE:	Completion Timeout Control Capability
DATE:	27 April 2005, updated 3 November 2005
AFFECTED DOCUMENT:	PCI Express Base Specification Revision 1.1
SPONSOR:	David Harriman, Intel Corporation

Part I

1 Summary of the Functional Changes

Adds a capability to disable the Completion Timeout mechanism and a capability to allow system firmware/software to set the Completion Timeout time value.

The disable capability is required for Endpoints and certain bridges, and optional for Root Ports.

The programmable timeout capability is optional. For this capability, the time ranges defined such that implementation is simplified by allowing a wide range of timeout values within each bin. These ranges are separated by illegal value ranges, so that it is possible for software to ensure that devices programmed to different settings (e.g. for hierarchical timeout) will time out at different times [as a counter example: If the ranges were contiguous, a device at the top end of one range would timeout at about the same time as a device operating at the lower end of the next bin].

2 Benefits as a Result of the Changes

Improved control of the Completion Timeout mechanism to reduce the likelihood of false triggering and enable greater situational appropriateness of Completion Timeout handling.

3 Assessment of the Impact

Some increase in hardware and software complexity to take advantage of the new capabilities provided.

4 Analysis of the Hardware Implications

The required disable mechanism has minimal hardware impact. The optional programmability mechanism allows implementation flexibility to provide more/less capability according to product requirements.

5 Analysis of the Software Implications

No impact to existing software. New software will optionally be able to use the new capabilities provided.

Part II

Detailed Description of the change

In Section 2.8:

...

PCI Express devices that issue Requests requiring Completions must implement the Completion Timeout mechanism. An exception is made for Configuration Requests (see below). The Completion Timeout mechanism is activated for each Request that requires one or more Completions when the Request is transmitted. Since PCI Express Switches do not autonomously initiate Requests that need Completions, the requirement for Completion Timeout support is limited only to Root Complexes, PCI Express-PCI Bridges, and Endpoint devices.

The Completion Timeout mechanism may be disabled by configuration software. The Completion Timeout limit is set in the Completion Timeout Value field of the Device Control 2 register. See Section <!!!ref>.

~~This specification defines the following range for the minimum/maximum acceptable timer values for the Completion Timeout mechanism:~~

- ~~The Completion Timeout timer must not expire (i.e., cause a timeout event) in less than 50 µs.~~

  ~~It is strongly recommend that unless an application requires this level of timer granularity the minimum time should not expire in less than 10 ms.~~

- ~~The Completion Timeout timer must expire if a Request is not completed in 50 ms.~~

A Completion Timeout is a reported error associated with the Requester device/function (see Section 6.2).

Note: A Memory Read Request for which there are multiple Completions must be considered completed only when all Completions have been received by the Requester. If some, but not all, requested data is returned before the Completion Timeout timer expires, the Requester is permitted to keep or to discard the data that was returned prior to timer expiration.

Completion timeouts for Configuration Requests have special requirements for the support of PCI Express to PCI/PCI-Express bridges. PCI Express to PCI/PCI-X Bridges, by default, are not enabled to return Configuration Request Retry Status (CRS) for Configuration Requests to a PCI/PCI-X device behind the Bridge. This may result in lengthy completion delays that must be comprehended by the Completion Timeout value in the Root Complex. System software may enable PCI Express to PCI/PCI-X Bridges to return Configuration Request Retry Status by setting the Bridge Configuration Retry Enable bit in the Device Control register, subject to the restrictions noted in the PCI Express to PCI/PCI-X Bridge Specification, Rev. 1.0.

...

Building upon the PCI Express Capability Structure Expansion ECR/ECN, in Device Capabilities 2 and Device Control 2 Register sections, change as shown:

7.8.15 Device Capabilities 2 Register (Offset 24h)

~~This section is a placeholder – There are no capabilities that require this register.~~

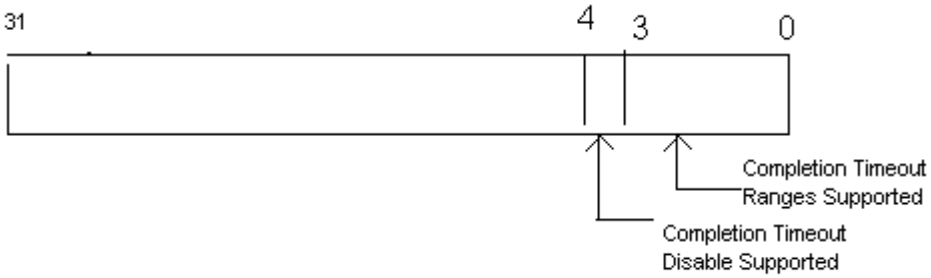

Figure 7-XX: Device Capabilities 2 Register

Table 7-XX: Device Capabilities 2 Register

Bit Location	Register Description	Attributes
3:0	**Completion Timeout Ranges Supported** – This field indicates device support for the optional Completion Timeout programmability mechanism. This mechanism allows system software to modify the Completion Timeout value. This field is applicable only to Root Ports, Endpoints that issue requests on their own behalf, and PCI Express to PCI/PCI-X Bridges that take ownership of requests issued on PCI Express. For all other devices this field is reserved and must be hardwired to 0000b. Four time value ranges are defined: Range A: 50us to 10ms Range B: 10ms to 250ms Range C: 250ms to 4s Range D: 4s to 64s Bits are set according to the table below to show timeout value ranges supported. 0000b Completion Timeout programming not supported – the device must implement a timeout value in the range 50us to 50ms. 0001b Range A 0010b Range B 0011b Ranges A & B 0110b Ranges B & C 0111b Ranges A, B & C 1110b Ranges B, C & D 1111b Ranges A, B, C & D All other values are reserved. It is strongly recommended that the Completion Timeout mechanism not expire in less than 10ms	HwInit

Bit Location	Register Description	Attributes
4	**Completion Timeout Disable Supported** – A value of 1b indicates support for the Completion Timeout Disable mechanism. The Completion Timeout Disable mechanism is required for Endpoints that issue requests on their own behalf and PCI Express to PCI/PCI-X Bridges that take ownership of requests issued on PCI Express. This mechanism is optional for Root Ports. For all other devices this field is reserved and must be hardwired to 0b.	RO

7.8.16 Device Control 2 Register (Offset 28h)

~~This section is a placeholder – There are no capabilities that require this register.~~

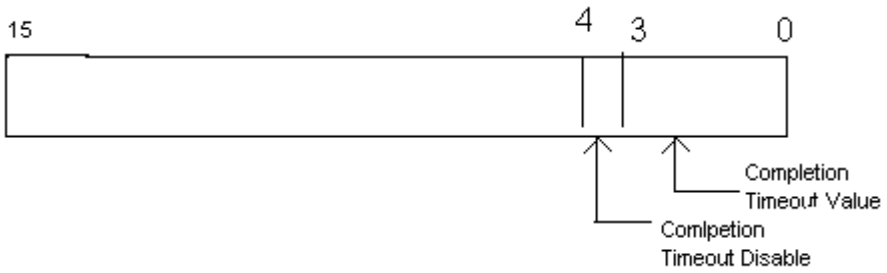

Figure 7-XX: Device Control 2 Register

Table 7-XX: Device Control 2 Register

Bit Location	Register Description	Attributes
3:0	**Completion Timeout Value** – In Devices that support Completion Timeout programmability, this field allows system software to modify the Completion Timeout value. This field is applicable to Root Ports, Endpoints that issue requests on their own behalf, and PCI Express to PCI/PCI-X Bridges that take ownership of requests issued on PCI Express. For all other devices this field is reserved and must be hardwired to 0000b. A Device that does not support this optional capability must hardwire this field to 0000b and is required to implement a timeout value in the range 50us to 50ms. Devices that support Completion Timeout programmability must support the values given below corresponding to the programmability ranges indicated in the Completion Timeout Values Supported field. Defined encodings: 0000b Default range: 50us to 50ms It is strongly recommended that the Completion Timeout mechanism not expire in less than 10ms Values available if Range A (50us to 10ms) programmability range is supported: 0001b 50us to 100us 0010b 1ms to 10ms Values available if Range B (10ms to 250ms) programmability range is supported: 0101b 16ms to 55ms 0110b 65ms to 210ms Values available if Range C (250ms to 4s) programmability range is supported: 1001b 260ms to 900ms 1010b 1s to 3.5s Values available if the Range D (4s to 64s) programmability range is supported: 1101b 4s to 13s 1110b 17s to 64s Values not defined above are Reserved. Software is permitted to change the value in this field at any time. For requests already pending when the Completion Timeout Value is changed, hardware is permitted to use either the new or the old value for the outstanding requests, and is permitted to base the start time for each request either on when this value was changed or on when each request was issued. The default value for this field is 0000b.	RW

Bit Location	Register Description	Attributes
4	**Completion Timeout Disable** – When set to 1b, this bit disables the Completion Timeout mechanism. This field is required for all devices that support the Completion Timeout Disable Capability. Software is permitted to set or clear this bit at any time. When set, the Completion Timeout detection mechanism is disabled. If there are outstanding requests when the bit is cleared, it is permitted but not required for hardware to apply the completion timeout mechanism to the outstanding requests. If this is done, it is permitted to base the start time for each request on either the time this bit was cleared or the time each request was issued. The default value for this bit is 0b.	RW